MADNESS

From Math to Peta-App

Robert J. Harrison

harrisonrj@ornl.gov

robert.harrison@utk.edu

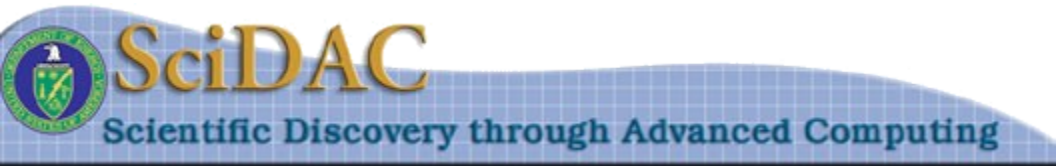

Mission of the ORNL National Leadership Computing Facility (NLCF)

- field the most powerful capability computers for scientific research
- select a few time sensitive problems of national importance that can take advantage of these systems
- join forces with the selected scientific teams to deliver breakthrough science.

Cray "Baker" – 1 PF System

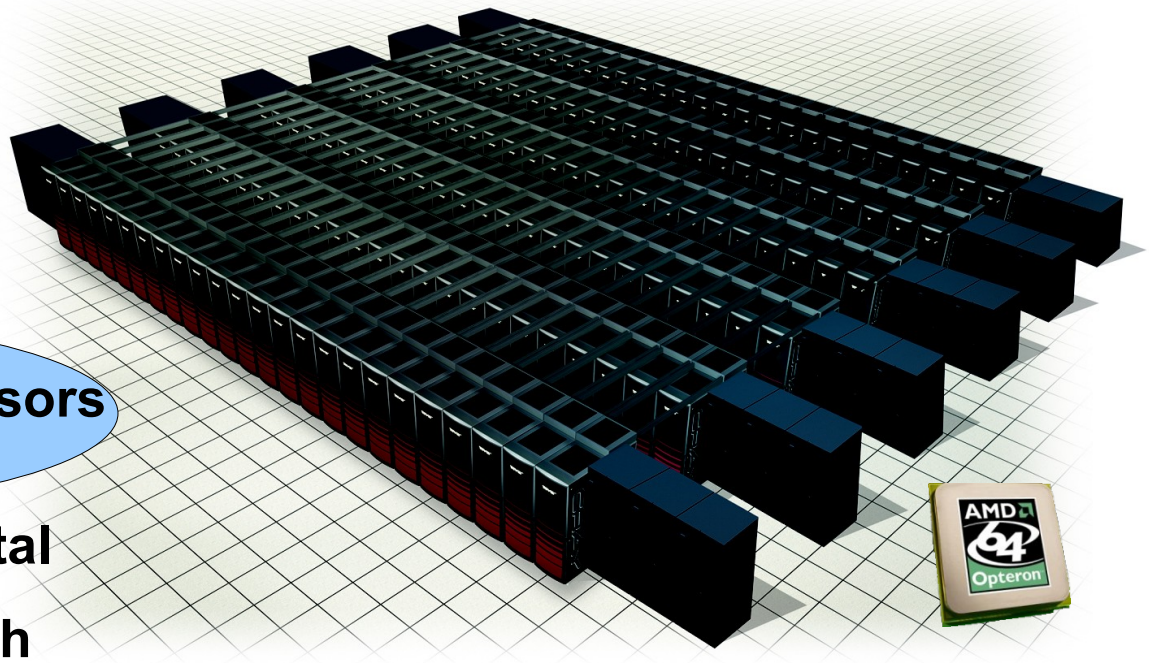

FY 2009: Cray "Baker"

- 1 Petaflops system
- 37 Gigaflops processor
- 27,888 quad-core processors Barcelona 2.3 GHz
- 2 GB per core; 223 TB total
- 200+ GB/s disk bandwidth
- 13,944 dual-socket 8-core SMP "nodes" with 16 GB
- 6.5 MW system power
- 150 Cabinets, 3,500 ft²

- Liquid cooled
- Compute node Linux operating system
- Torus interconnect

Now beginning to work!

111,552 cores @ 9.2GFlop/s

Full details to be announced at SC08

Univ. of Tennessee & ORNL Partnership
National Institute for Computational Sciences

- UT is building a new NSF supercomputer center from the ground up
  - Building on strengths of UT and ORNL
  - Operational in May 2008
- Series of computers culminating in a 1 PF system in 2009
  - Initial delivery (May 2008)
    - 4512 quad-core Opteron processors (170 TF)
  - Cray "Baker" (2009)
    - Multi-core Opteron processors; 100 TB; 2.3 PB of disk space

O(1) programmers … O(10,000) nodes … O(100,000) processors … O(10,000,000) threads

- Complexity kills … sequential or parallel
- Expressing/managing concurrency at the petascale
  - It is too trite to say that the parallelism is in the physics
  - Must express and discover parallelism at more levels
  - Low level tools (MPI, Co-Array Fortran, UPC, …) don't discover parallelism or hide complexity or facilitate abstraction
- Management of the memory hierarchy
  - Memory will be deeper ; less uniformity between vendors
  - Need tools to automate and manage this, even at runtime

The way forward demands a change in paradigm

- by us chemists, the funding agencies, and the supercomputer centers

- A communal effort recognizing the increased cost and complexity of code development for modern theory at the petascale

- Re-emphasizing basic and advanced theory and computational skills in undergraduate and graduate education

Computational Chemistry Endstation

International collaboration spanning 7 universities and 6 national labs

- Led out of UT/ORNL
- Focus
  - Actinides, Aerosols, Catalysis
- ORNL Cray XT, ANL BG/L

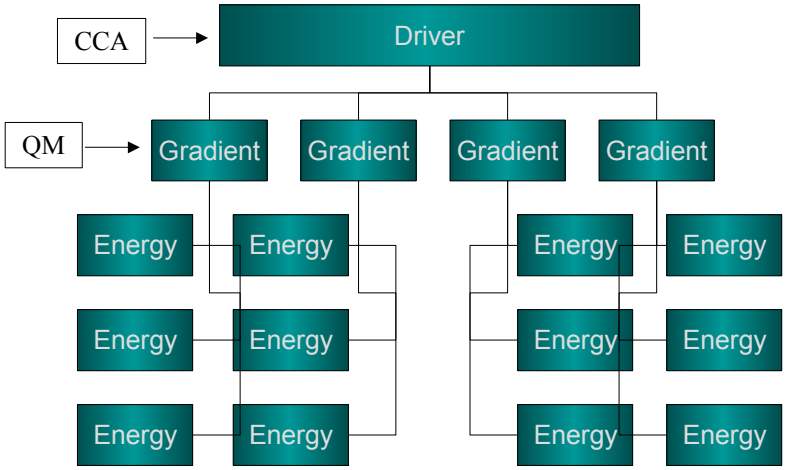

TL Windus

Capabilties:

- Chemically accurate thermochemistry
  - Many-body methods required
- Mixed QM/QM/MM dynamics
  - Accurate free-energy integration
  - Simulation of extended interfaces
- Families of relativistic methods

Participants:

- Harrison, UT/ORNL
- Sherrill, GATech
- Gordon, Windus, Iowa State / Ames
- Head-Gordon, U.C. Berkeley / LBL
- Crawford, Valeev, VTech.
- Bernholc, NCSU
- (Knowles, U. Cardiff, UK)
- (de Jong, PNNL)
- (Shepard, ANL)
- (Sherwood, Daresbury, UK)

Linear/Reduced Scaling Methods

- Non-linear scaling of the computational cost is not acceptable for massively parallel software
  - E.g., if cost = $O(N^3)$ then a computer that 1000x faster can only run a calculation 10x larger
- Must work on all of
  - Theory
  - Numerical representation
  - Algorithm
  - Efficient implementation

MADNESS

Multiresolution
Adaptive
Numerical
Scientific
Simulation

Multiresolution Adaptive Numerical Scientific Simulation

Ariana Beste[1], George I. Fann[1], Robert J. Harrison[1,2], Rebecca Hartman-Baker[1], Jun Jia[1], Shinichiro Sugiki[1]
[1]Oak Ridge National Laboratory, [2]University of Tennessee, Knoxville

Gregory Beylkin[4], Fernando Perez[4], Lucas Monzon[4], Martin Mohlenkamp[5] and others
[4]University of Colorado, [5]Ohio University

Hideo Sekino[6] and Takeshi Yanai[7]
[6]Toyohashi University of Technology, [7]Institute for Molecular Science, Okazaki

harrisonrj@ornl.gov

SciDAC
Scientific Discovery through Advanced Computing

UT-BATTELLE
Oak Ridge National Laboratory

The DOE funding

- This work is funded by the U.S. Department of Energy, the divisions of Advanced Scientific Computing Research and Basic Energy Science, Office of Science, under contract DE-AC05-00OR22725 with Oak Ridge National Laboratory. This research was performed in part using
  - resources of the National Energy Scientific Computing Center which is supported by the Office of Energy Research of the U.S. Department of Energy under contract DE-AC03-76SF0098,
  - and the Center for Computational Sciences at Oak Ridge National Laboratory under contract DE-AC05-00OR22725 .

SciDAC

Scientific Discovery through Advanced Computing

Multiresolution chemistry objectives

- Scaling to 1+M processors ASAP
- Complete elimination of the basis error
  - One-electron models (e.g., HF, DFT)
  - Pair models (e.g., MP2, CCSD, …)
- Correct scaling of cost with system size
- General approach
  - Readily accessible by students and researchers
  - Higher level of composition
  - Direct computation of chemical energy differences
- New computational approaches
  - *Fast algorithms with guaranteed precision*

The mathematicians …

Gregory Beylkin
http://amath.colorado.edu/faculty/beylkin/

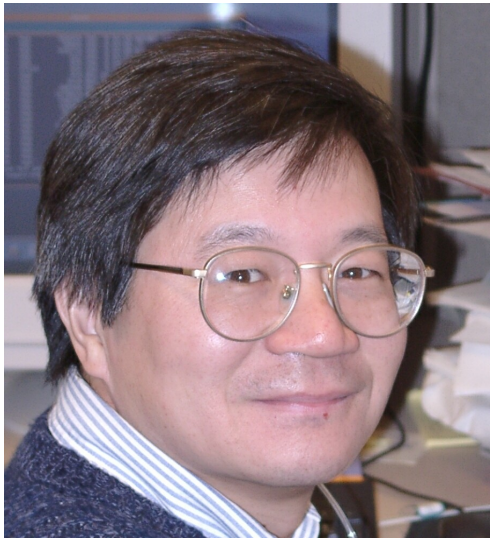

George I. Fann
fanngi@ornl.gov

Molecular orbitals of water

2-d contour plot

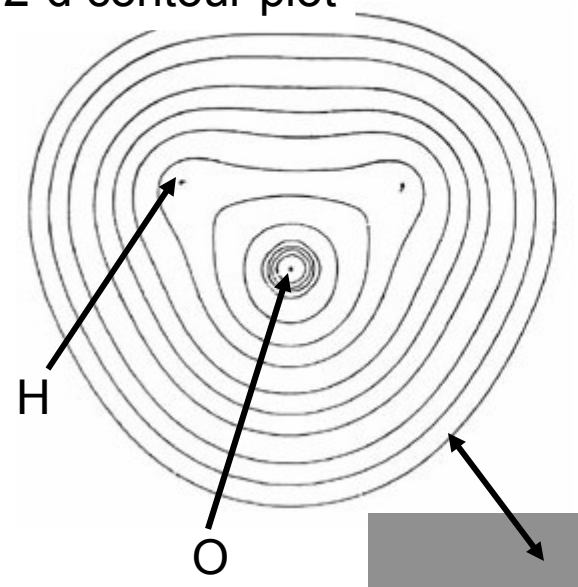

Iso-surfaces are 3-d contour plots – they show the surface upon which the function has a particular value

Water has 10 electrons (8 from oxygen, 1 from each hydrogen).

It is closed-shell, so it has 5 molecular orbitals each occupied with two electrons.

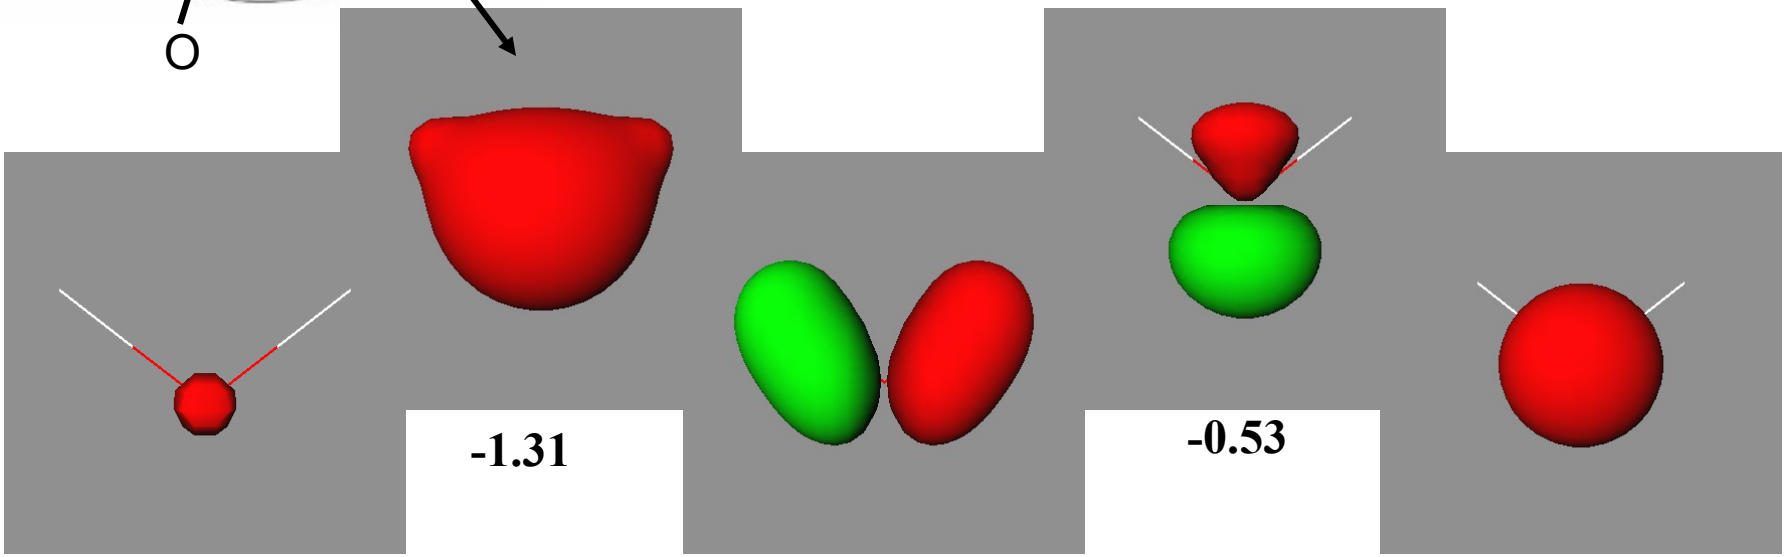

The energy of each orbital in atomic units

Linear Combination of Atomic Orbitals (LCAO)

- Molecules are composed of (weakly) perturbed atoms
  - Use finite set of atomic wave functions as the basis
  - Hydrogen-like wave functions are exponentials
- E.g., hydrogen molecul

$$1s(r) = \mathrm{e}^{-|r|}$$

$$\phi(r) = \mathrm{e}^{-|r-a|} + \mathrm{e}^{-|r-b|}$$

- Smooth function of molecular geometry
- MOs: cusp at nucleus with exponential decay

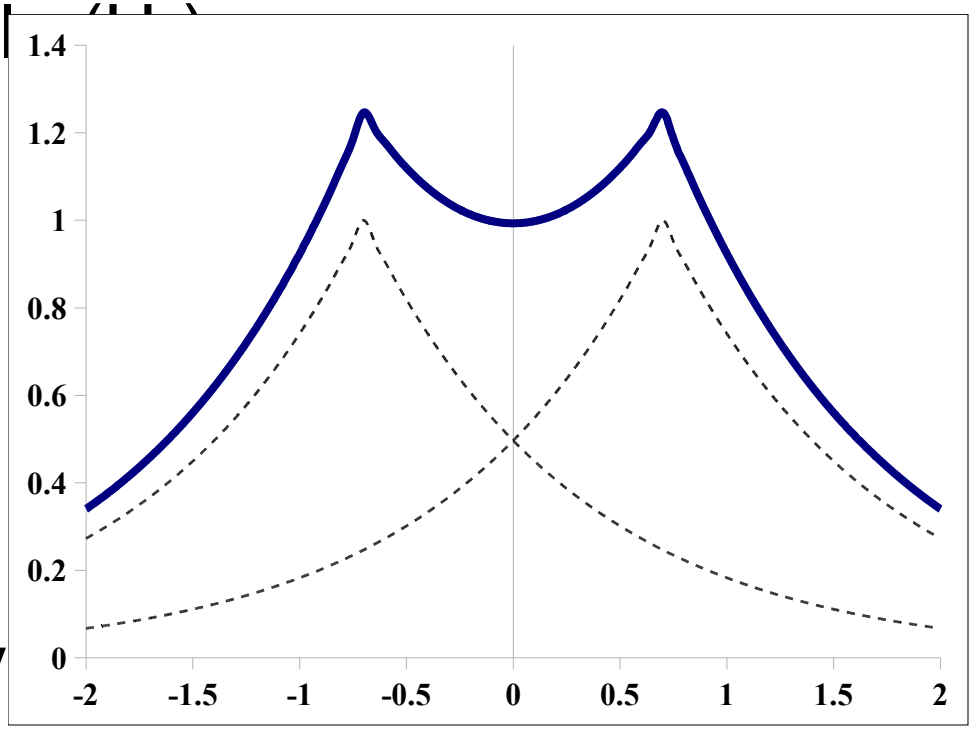

LCAO with Gaussian Functions

- Cannot compute integrals over exponential orbitals
- Boys (1950) noted that Gaussians are feasible
  - 6D integral reduced to 1D integrals which are tabulated once and stored (related to error function)
- Gaussian functions form a complete basis
  - With enough terms any radial function can be approximated to any precision using a linear combination of Gaussian functions

$$f(r) = \sum_{i=1}^{N} c_i e^{-a_i r^2} + O(\epsilon)$$

LCAO

- A fantastic success, but …
- Basis functions have extended support
  - causes great inefficiency in high accuracy calculations (functions on different centers overlap)
  - origin of non-physical density matrix
- Basis set superposition error (BSSE)
  - incomplete basis on each center leads to over-binding as atoms are brought together
- Linear dependence problems
  - accurate calculations require balanced approach to a complete basis on every atom
  - molecular basis can have severe linear dependence
- Must extrapolate to complete basis limit
  - unsatisfactory and not feasible for large systems

Essential techniques for fast computation

- Multiresolution

$$V_0 \subset V_1 \subset \cdots \subset V_n$$

$$V_n = V_0 + \left(V_1 - V_0\right) + \cdots + \left(V_n - V_{n-1}\right)$$

- Low-separation rank

$$f(x_1, \ldots, x_n) = \sum_{l=1}^{M} \sigma_l \prod_{i=1}^{d} f_i^{(l)}(x_i) + O(\epsilon)$$

$$\left\| f_i^{(l)} \right\|_2 = 1 \qquad \sigma_l > 0$$

- Low-operator rank

$$A = \sum_{\mu=1}^{r} u_\mu \sigma_\mu v_\mu^T + O(\epsilon)$$

$$\sigma_\mu > 0 \qquad v_\mu^T v_\lambda = u_\mu^T u_\lambda = \delta_{\mu\nu}$$

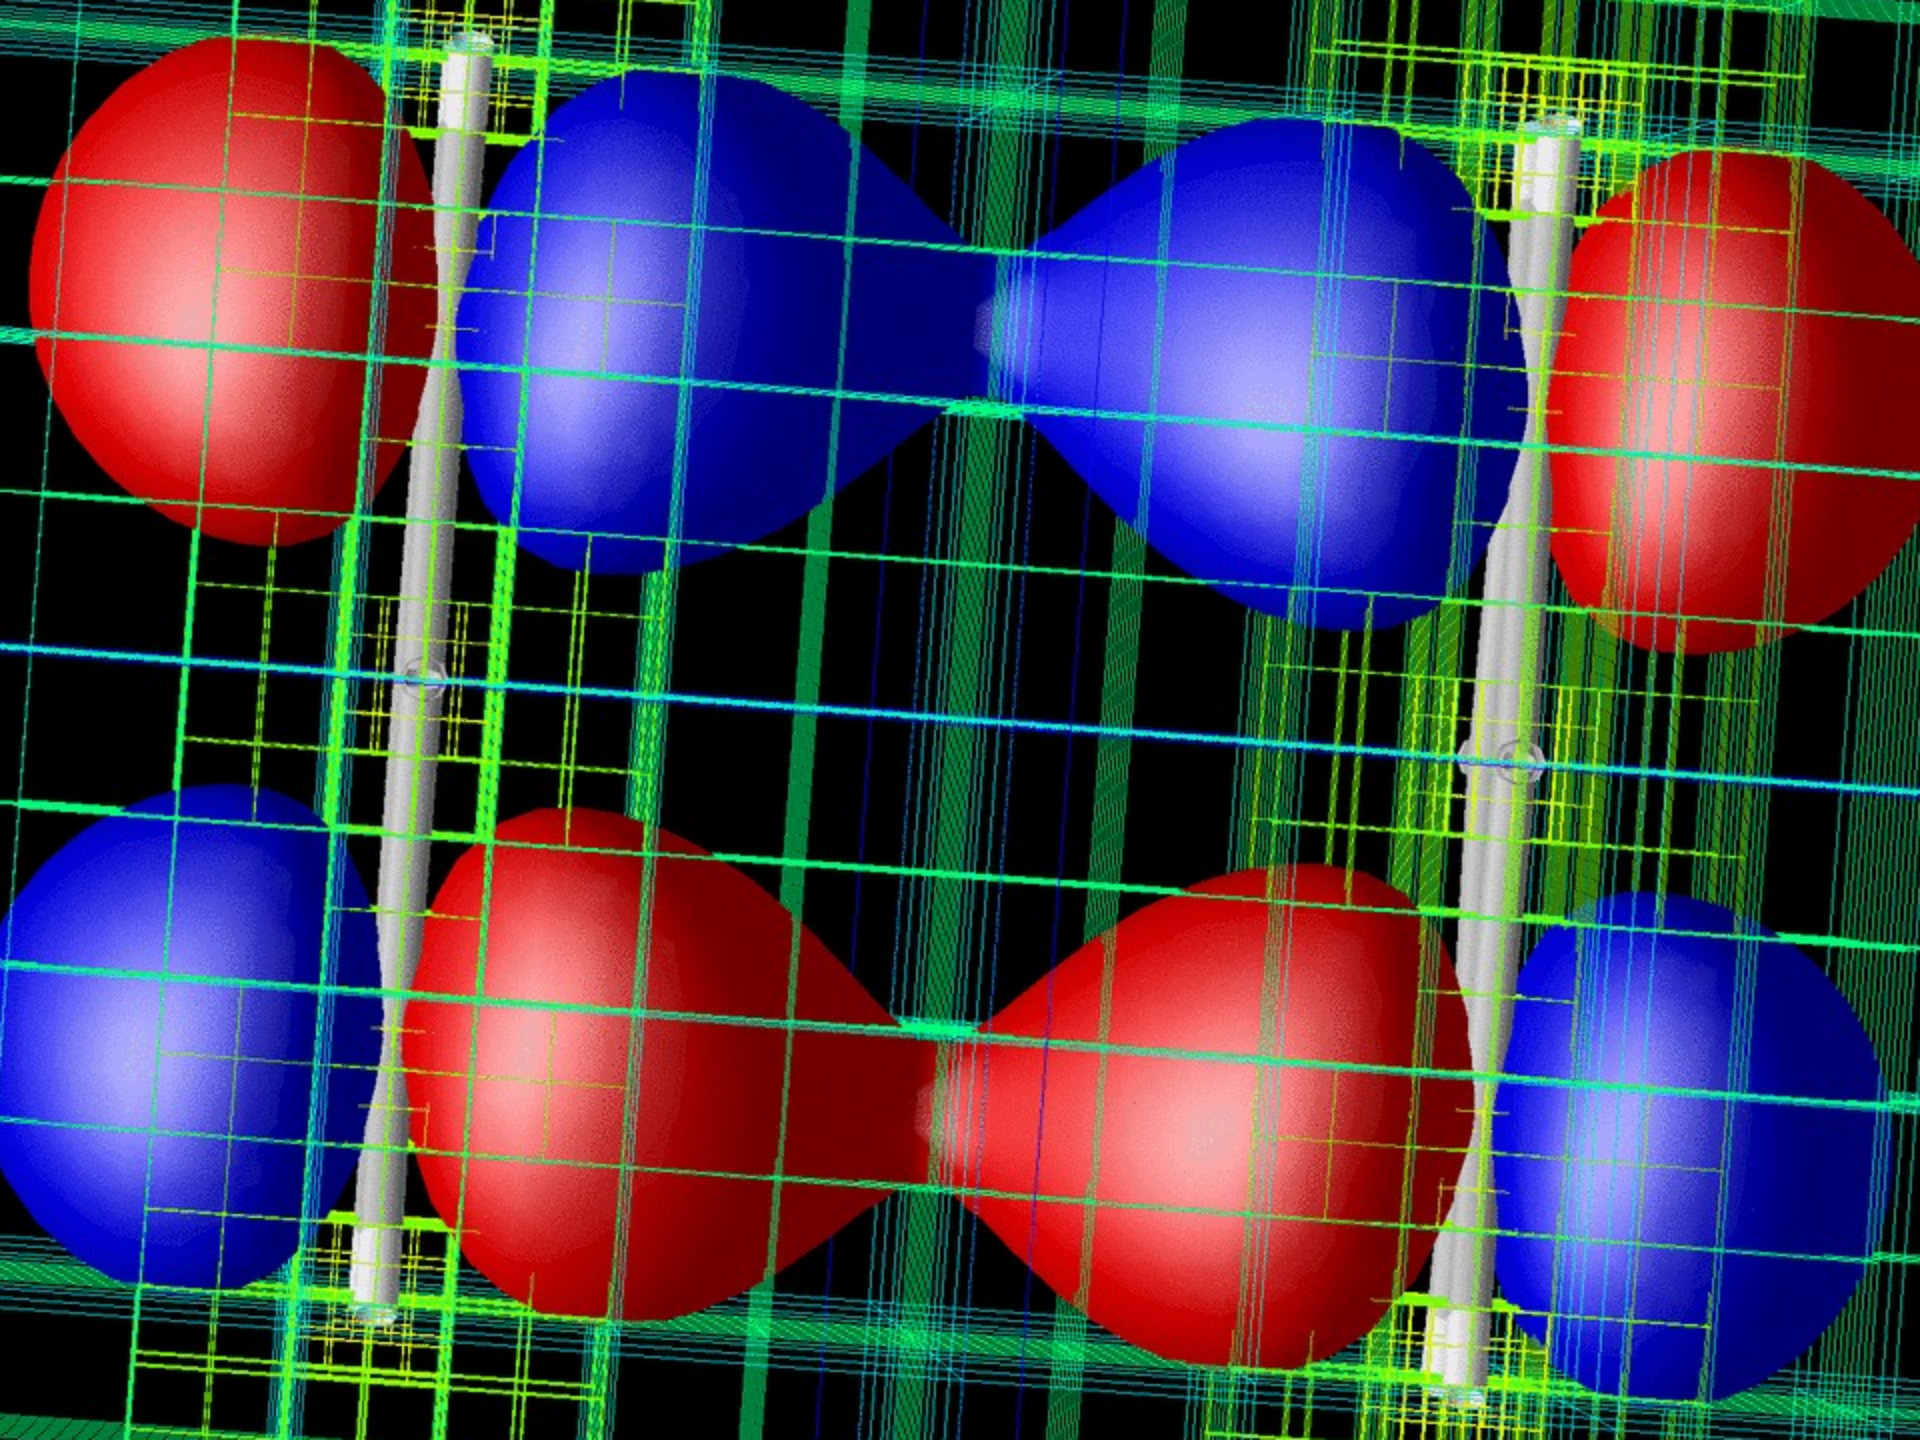

Please forget about wavelets

- They are not central
- Wavelets are a convenient basis for spanning V_n-V_{n-1} and understanding its properties
- But you don't actually need to use them
  - MADNESS does still compute wavelet coefficients, but *Beylkin's new code does not*
- Please remember this …
  - Discontinuous spectral element with multi-resolution and separated representations for fast computation with guaranteed precision in many dimensions.

Computational kernels

- Discontinuous spectral element
  - In each “box” a tensor product of coefficients
  - Most operations are small matrix-multiplication

$$r_{i'j'k'} = \sum_{ijk} s_{ijk} c_{ii'} c_{jj'} c_{kk'} = \sum_{k} \left(\sum_{j} \left(\sum_{i} s_{ijk} c_{ii'} \right) c_{jj'} \right) c_{kk'}$$

$$\Rightarrow r = \left(\left(s^T c \right)^T c \right)^T c$$

  - Typical matrix dimensions are 2 to 30
  - E.g., $(20,400)^T$ * (20,20)
  - Often low rank

Speed relative to MKL, Goto, ATLAS on Intel Xeon 5355 for $(20,400)^T*(20,n)$.

n	MKL	Goto	ATLAS		n	MKL	Goto	ATLAS
2	6.25	4.1667	5		16	0.8966	1.2581	2.0708
4	3.1042	3.6341	4.6563		18	1.7763	1.3636	2.4545
6	4.375	2.625	5.122		20	0.9556	1.2727	2.6168
8	1.3132	2.0427	5.1957		22	1.6416	1.2968	2.7308
10	2.7368	1.9549	5.3061		24	0.9638	1.2208	1.9664
12	1.0605	1.5843	2.4352		26	1.5337	1.2814	2.1295
14	2.0323	1.4737	2.1356		28	0.8411	1.0588	2.0301

Ratio of flops/cycle of MTXMQ to listed code: value > 1 indicates we are faster

XT5 single core FLOPs/cycle

(nj, ni)T*(nj,nk)				
ni	nj	nk	MTXMQ	ACML
400	2	20	2.55	0.95
400	4	20	2.62	1.50
400	6	20	2.60	1.79
400	8	20	2.56	2.02
400	10	20	2.58	2.12
400	12	20	2.64	2.27
400	14	20	2.90	2.35
400	16	20	2.80	2.46
400	18	20	2.74	2.49
400	20	20	2.89	2.58

nested transform (nj, ni)T*(nj,nk)				
ni	nj	nk	MTXMQ	ACML
4	2	2	0.10	0.07
16	4	4	1.04	0.51
36	6	6	1.74	0.99
64	8	8	2.33	1.56
100	10	10	2.61	1.80
144	12	12	2.69	2.12
196	14	14	2.94	2.17
256	16	16	2.97	2.41
324	18	18	2.93	2.38
400	20	20	3.03	2.49
484	22	22	3.01	2.52
576	24	24	3.09	2.73
676	26	26	3.02	2.73
784	28	28	2.87	2.87
900	30	30	2.88	2.81

L2 cache is 512Kb = 2*32^3 doubles
- hence expect good multi-core scaling
- don’t have actual data ... yet.

Applications under active development

- DFT & HF for electrons
  - Energies, gradients, spectra, non-linear optical properties, Raman intensities (Harrison, Sekino, Yanai)
  - Molecules & periodic systems (Eguilez and Thornton)
- Atomic and molecular physics
  - Exact dynamics of few electron systems in strong fields (Krstic and Vence), MCSCF for larger systems
- Nuclear structure
  - G. Fann, et al.
- Preliminary studies in fusion and climate

Path to linear scaling HF & DFT

- Need speed and precision
  - Absolute error cost $O\left(N \ln N / \epsilon\right)$
  - Relative error cost $O\left(N \ln 1 / \epsilon\right)$
- Coulomb potential
- HF exchange potential
- Orbital update
- Orthogonalization and or diagonalization
- Linear response properties

High-level composition

- Close to the physics

$$E = \langle \psi | -\frac{1}{2}\nabla^2 + V | \psi \rangle + \int \psi^2(x) \frac{1}{|x-y|} \psi^2(y)\, dx\, dy$$

```
operatorT op = CoulombOperator(k, rlo, thresh);
functionT rho = psi*psi;
double twoe = inner(apply(op,rho),rho);
double pe = 2.0*inner(Vnuc*psi,psi);
double ke = 0.0;
for (int axis=0; axis<3; axis++) {
    functionT dpsi = diff(psi,axis);
    ke += inner(dpsi,dpsi);
}
double energy = ke + pe + twoe;
```

High-level composition

- Express ***ALL*** available parallelism without burdening programmer
  - Internally, MADNESS is looking after data and placement and scheduling of operations on individual functions
  - Programmer must express parallelism over multiple functions and operators
    - But is *not* responsible for scheduling or placement

High-level composition

- E.g., make the matrix of KE operator
  - All scalar operations include optional fence
    - E.g., functionT scale(const functionT& f, T scale, bool fence=true)
  - Internally, operations on vectors schedule all tasks with only one fence

$$\langle \phi_i | -\frac{1}{2}\nabla^2 | \phi_j \rangle = \frac{1}{2}\langle \nabla^T \phi_i \nabla \phi_j \rangle$$

```
Tensor<double>
kinetic_energy_matrix(World& world,
                      const vector<functionT>& v) {
  int n = v.size();
  Tensor<double> r(n,n);
  for (int axis=0; axis<3; axis++) {
     vector<functionT> dv = diff(world,v,axis);
     r += inner(world, dv, dv);
  }
  return r.scale(0.5);
}
```

MADNESS architecture

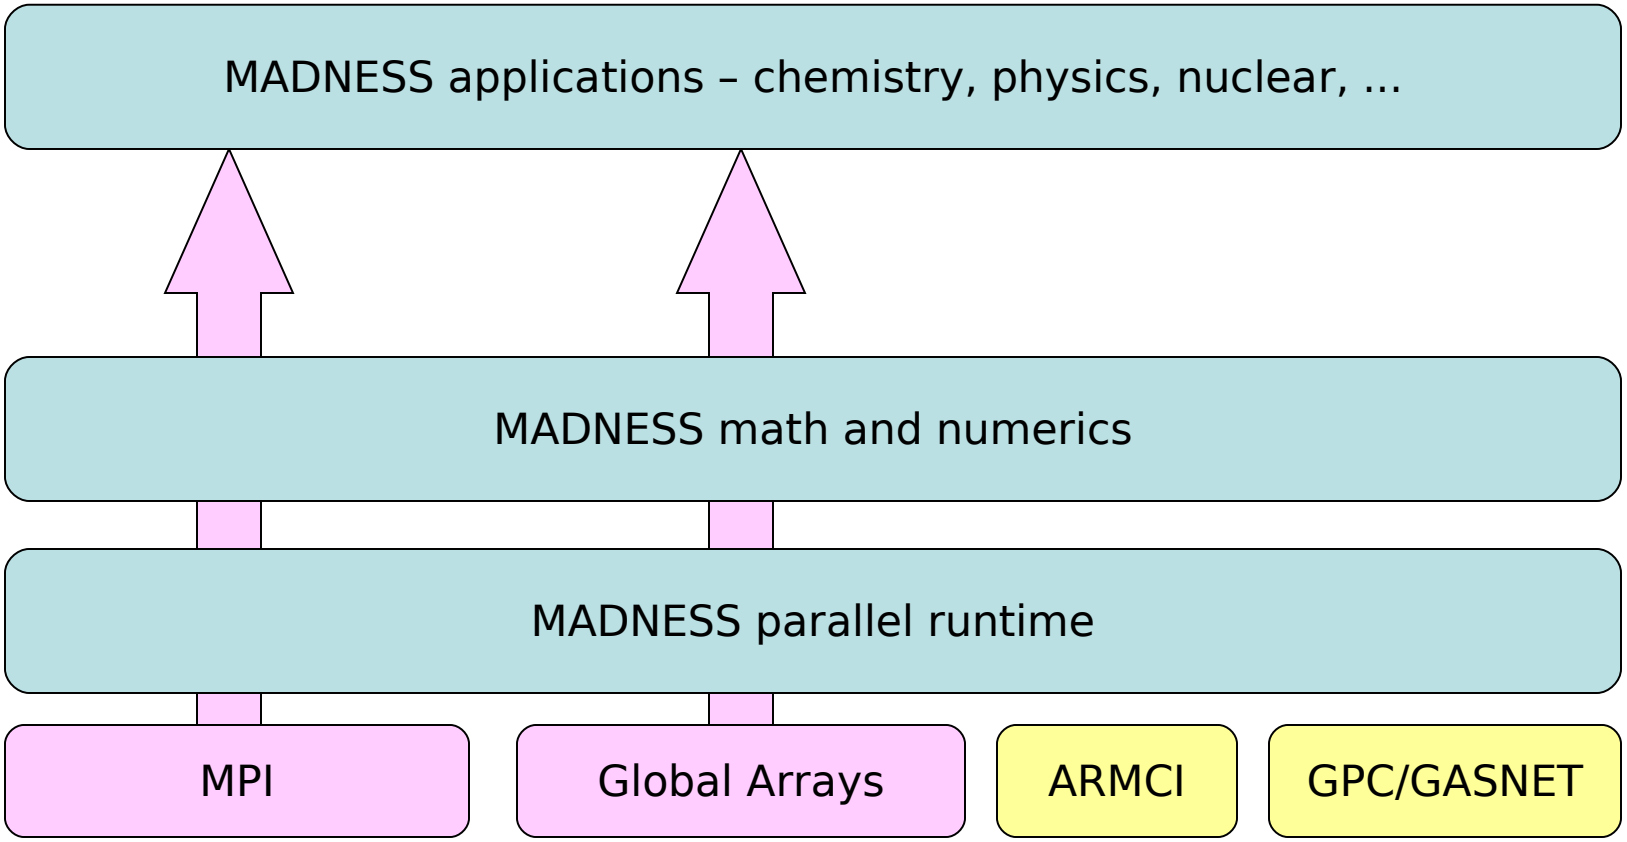

Intel Thread Building Blocks being considered for multicore

Runtime Objectives

- Scalability to 1+M processors ASAP
- Runtime responsible for
  - scheduling and placement,
  - managing data dependencies,
  - hiding latency, and
  - Medium to coarse grain concurrency
- Compatible with existing models
  - MPI, Global Arrays
- Borrow successful concepts from Cilk, Charm++, Python
- Anticipating next gen. languages

Key elements

- Futures for hiding latency and automating dependency management
- Global names and name spaces
- Non-process centric computing
  - One-sided messaging between objects
  - Retain place=process for MPI/GA legacy
- Dynamic load balancing
  - Data redistribution, work stealing, randomization

Multi-threaded architecture

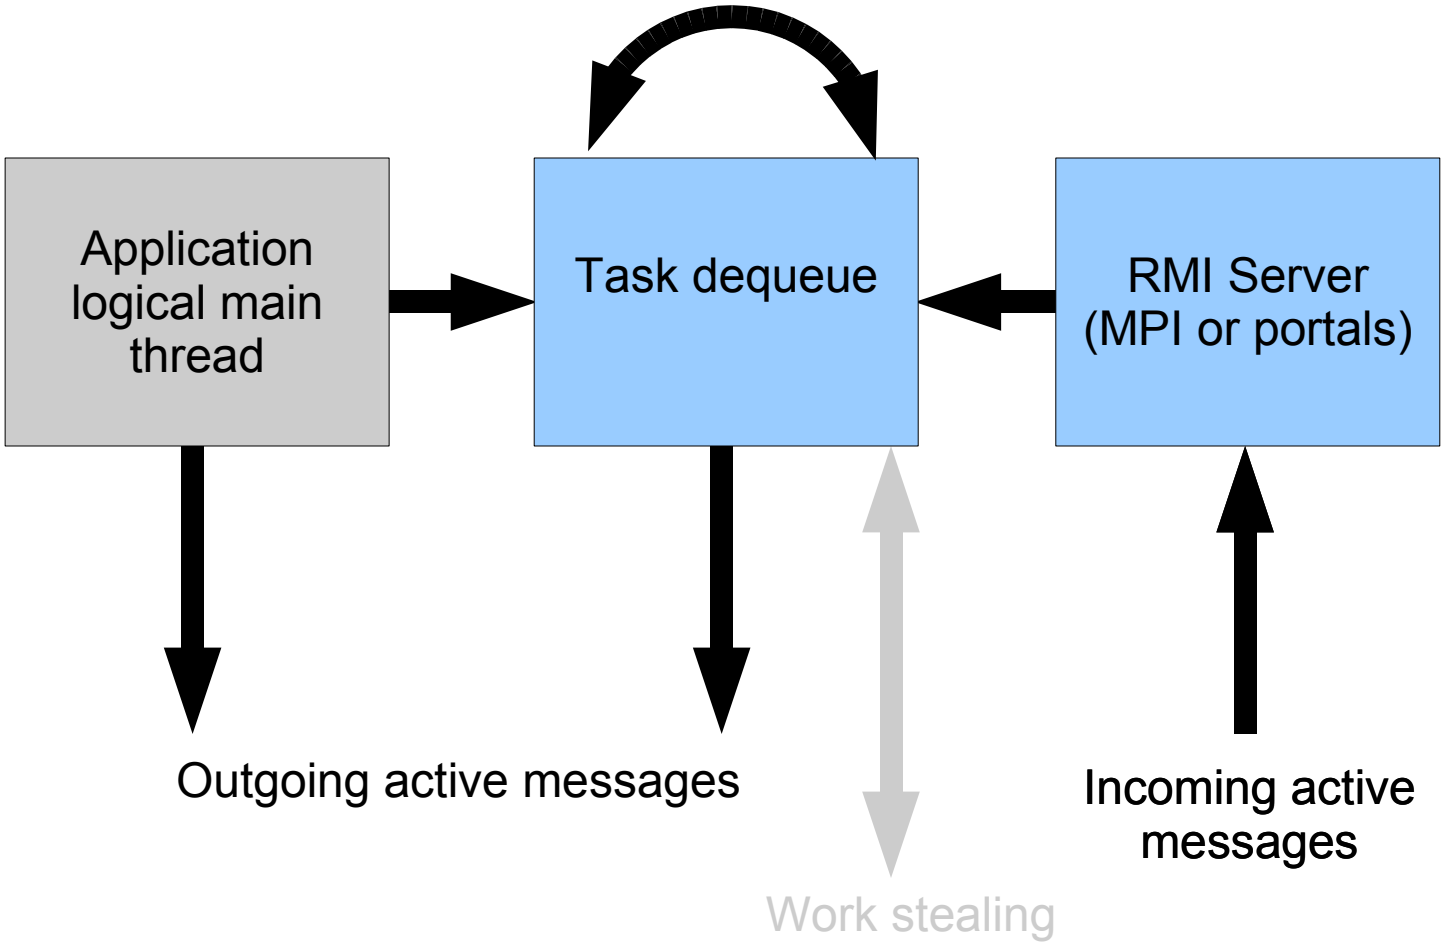

Issues

- Manual generation of continuations or closures
  - Tedious and error prone
    - Need compiler support
    - User-space threads/fibers can help (c.f., Cilk, Charm++)
- Transitioning between cache-oblivious and cache-aware algorithms
  - Essential for peak performance
- Hierarchical task expression
  - Better use of memory hierarchy
  - Throttle parallelism; enable DAG-based scheduling

Futures

- Result of an asynchronous computation
  - Cilk, Java, HPCLs
- Hide latency due to communication or computation
- Management of dependencies
  - Via callbacks

```
int f(int arg);
ProcessId me, p;

Future<int> r0=task(p, f, 0);
Future<int> r1=task(me, f, r0);

// Work until need result

cout << r0 << r1 << endl;
```

Process “me” spawns a new task in process “p” to execute `f(0)` with the result eventually returned as the value of future `r0`. This is used as the argument of a second task whose execution is deferred until its argument is assigned. Tasks and futures can register multiple local or remote callbacks to express complex and dynamic dependencies.

Virtualization of data and tasks

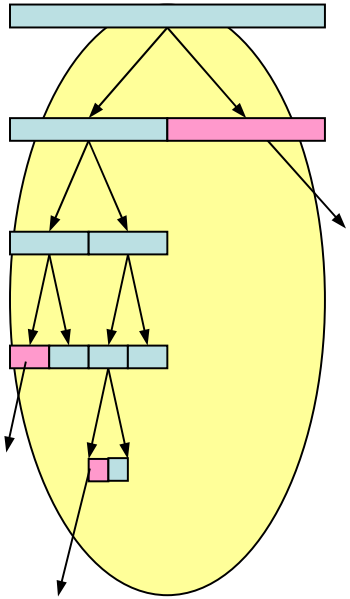

Parameter:
- MPI rank
- probe()
- set()
- get()

Task:
- Input parameters
- Output parameters
- probe()
- run()

```
Future Compress(tree):
      Future left = Compress(tree.left)
      Future right = Compress(tree.right)
      return Task(Op, left, right)

Compress(tree)
Wait for all tasks to complete
```

Benefits: Communication latency & transfer time largely hidden
Much simpler composition than explicit message passing
Positions code to use "intelligent" runtimes with work stealing
Positions code for efficient use of multi-core chips

```
#define WORLD_INSTANTIATE_STATIC_TEMPLATES
#include <world/world.h>
using namespace madness;
class Foo : public WorldObject<Foo> {
    const int bar;
public:
    Foo(World& world, int bar) : WorldObject<Foo>(world), bar(bar)
                {process_pending();}

    int get() const {return bar;}
};
int main(int argc, char** argv) {
    MPI::Init(argc, argv);
    madness::World world(MPI::COMM_WORLD);

    Foo a(world,world.rank()), b(world,world.rank()*10)

    for (ProcessID p=0; p<world.size(); p++) {
        Future<int> futa = a.send(p,&Foo::get);
        Future<int> futb = b.send(p,&Foo::get);
        // Could work here until the results are available
        MADNESS_ASSERT(futa.get() == p);
        MADNESS_ASSERT(futb.get() == p*10);
    }
    world.gop.fence();
    if (world.rank() == 0) print("OK!");
    MPI::Finalize();
}
```

Figure 1: Simple client-server program implemented using WorldObject.

Global Namespaces

- Specialize global names to containers
  - Hash table done
  - Arrays, etc., planned
- Replace global pointer (process+local pointer) with more powerful concept
-
- User definable map from keys to “owner” process

```
class Index;  // Hashable
class Value {
   double f(int);
};

WorldContainer<Index,Value> c;
Index i,j;  Value v;
c.insert(i,v);
Future<double> r =
   c.task(j,&Value::f,666);
```

A container is created mapping indices to values.

A value is inserted into the container.

A task is spawned in the process owning key `j` to invoke `c[j].f(666)`.

Namespaces are a large part of the elegance of Python and success of Charm++ (chares+arrays)

Abstraction Overheads

- If you are careful you win
  - *Increased performance and productivity*
  - This is the lesson of Global Arrays, Charm++, …
- Creating, executing, reaping a local, null task – 350ns (100K tasks, 3GHz Core2, Pathscale 3.0, -Ofast) dominated by new/delete
- Chain of 100K dependent tasks with the result of a task as the unevaluated argument of the previous task
  - ~1 us per task
- Creating a remote task adds overhead of inter-process communication which is on the scale of 5us (Cray XT).
  - Aggregation can reduce this.
- Switching between user-space threads <20ns

Summary

- Huge computational resources are rushing towards us
  - Tremendous scientific potential
  - Tremendous challenges
    - Research
    - Education
    - Community
- UT and ORNL are at the very center
  - Think of us when you have/want something fun and challenging to do

Extra Slides

HF Exchange (T. Yanai)

- HF or exact exchange
  - Features in the most successful XC functionals

$$\hat{K}\, f(x) = \sum_{i}^{\text{occupied}} n_i \phi_i(x) \int dy \frac{\phi_i(y) f(y)}{|x-y|}$$

  - Invariant to unitary rotation of occupied states with same occupation number
  - Localize the orbitals – only O(1) products but potential is still global
  - Compute potential only where orbital non-zero
  - Cost to apply to all orbitals circa O(N)

Orbital update

- Directly solve for localized orbitals that span space of occupied eigenfunctions
  - Rigorous error control from MRA refinement
  - Never construct the eigenfunctions
  - Update only diagonal multipliers
    - Off diagonal from localization process

$$\phi_i(x) = -(\hat{T} - \zeta)^{-1} \left((V + \zeta)\phi_i - \sum_j^{\text{occupied}} \phi_j(x)\epsilon_{ji} \right)$$

Inner products

- The most expensive term for plane wave codes leading to cost O(N^2 M)
- Inexpensive in MRA basis

$$\langle f|g\rangle = s_f^{00} \cdot s_g^{00} + \sum_{n=0}^{2^n-1} \sum_{l=0} d_f^{nl} \cdot d_g^{nl}$$

  - Orthogonal basis from local adaptive refinement implies zero/reduced work if
    - Functions do not overlap
    - Functions locally live at different length scales

Advancing Scientific Discovery

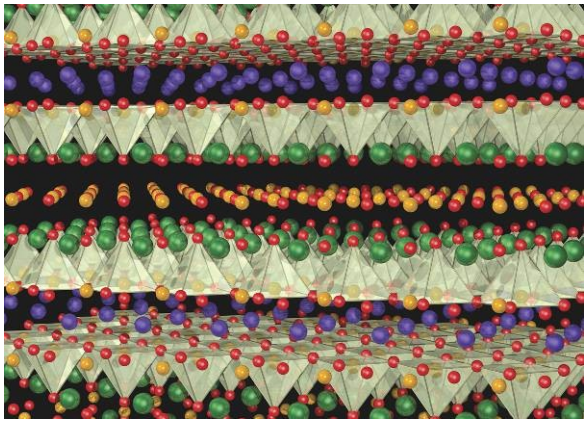

Resolved decades-long controversy about validity of 2D Hubbard model in predicting behavior of high-temperature superconducting cuprate planes

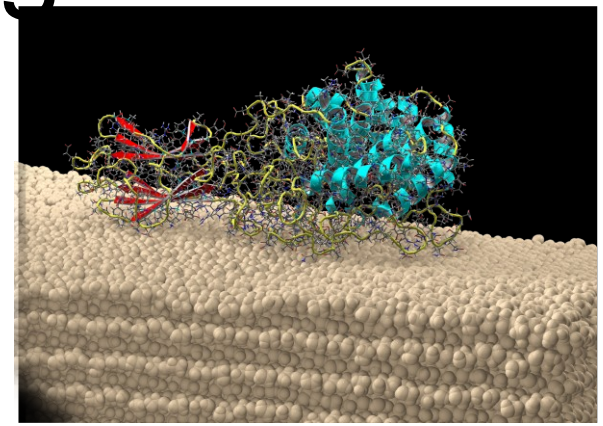

300K-atom models of cellulase enzyme on cellulose substrate reveal interior enzyme vibrations that influence reaction rates converting cellulose to ethanol

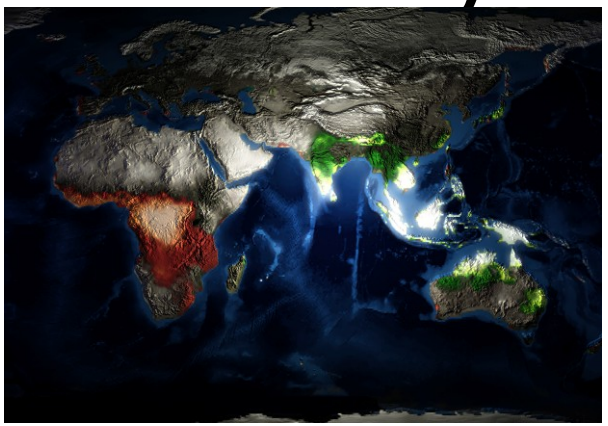

Addition and intercomparison of carbon-land models in new climate model is resolving key processes for carbon sources & sinks

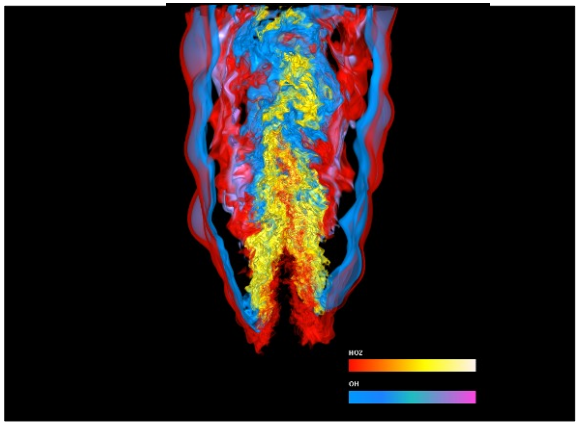

Turbulence chemistry revealed in study of lifted turbulent H_2/air jet flames in ignitive coflow relevant to diesel engines and gas turbines

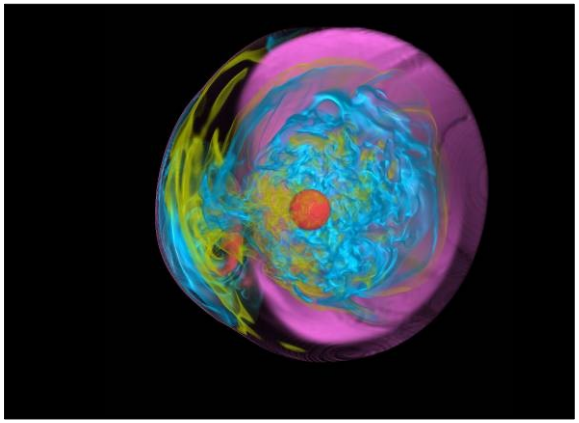

Instability of supernova shocks was discovered directly through simulation and core collapse pulsar mechanism was explained

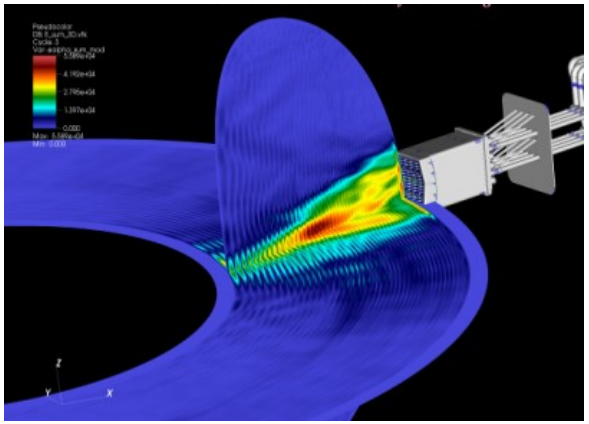

Providing increasing assurance that RF power will effectively heat ITER

Electron correlation

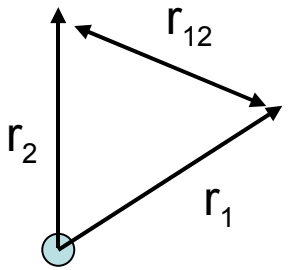

- All defects in the mean-field model are ascribed to electron correlation
- Consideration of singularities in the Hamiltonian imply that for a two-electron singlet atom (e.g., He)

$$\Psi(r_1, r_2, r_{12}) = 1 + \frac{1}{2} r_{12} + O(r_{12}^2) \quad \text{as} \quad r_{12} \to 0$$

- Include the inter-electron distance in the wavefunction
  - E.g., Hylleraas 1938 wavefunction for He

$$\Psi(r_1, r_2, r_{12}) = e^{-\xi(r_1 + r_2)} \left(1 + a\, r_{12} + \cdots\right)$$

  - Potentially very accurate, but not systematically improvable, and (until recently) not computationally feasible for many-electron systems

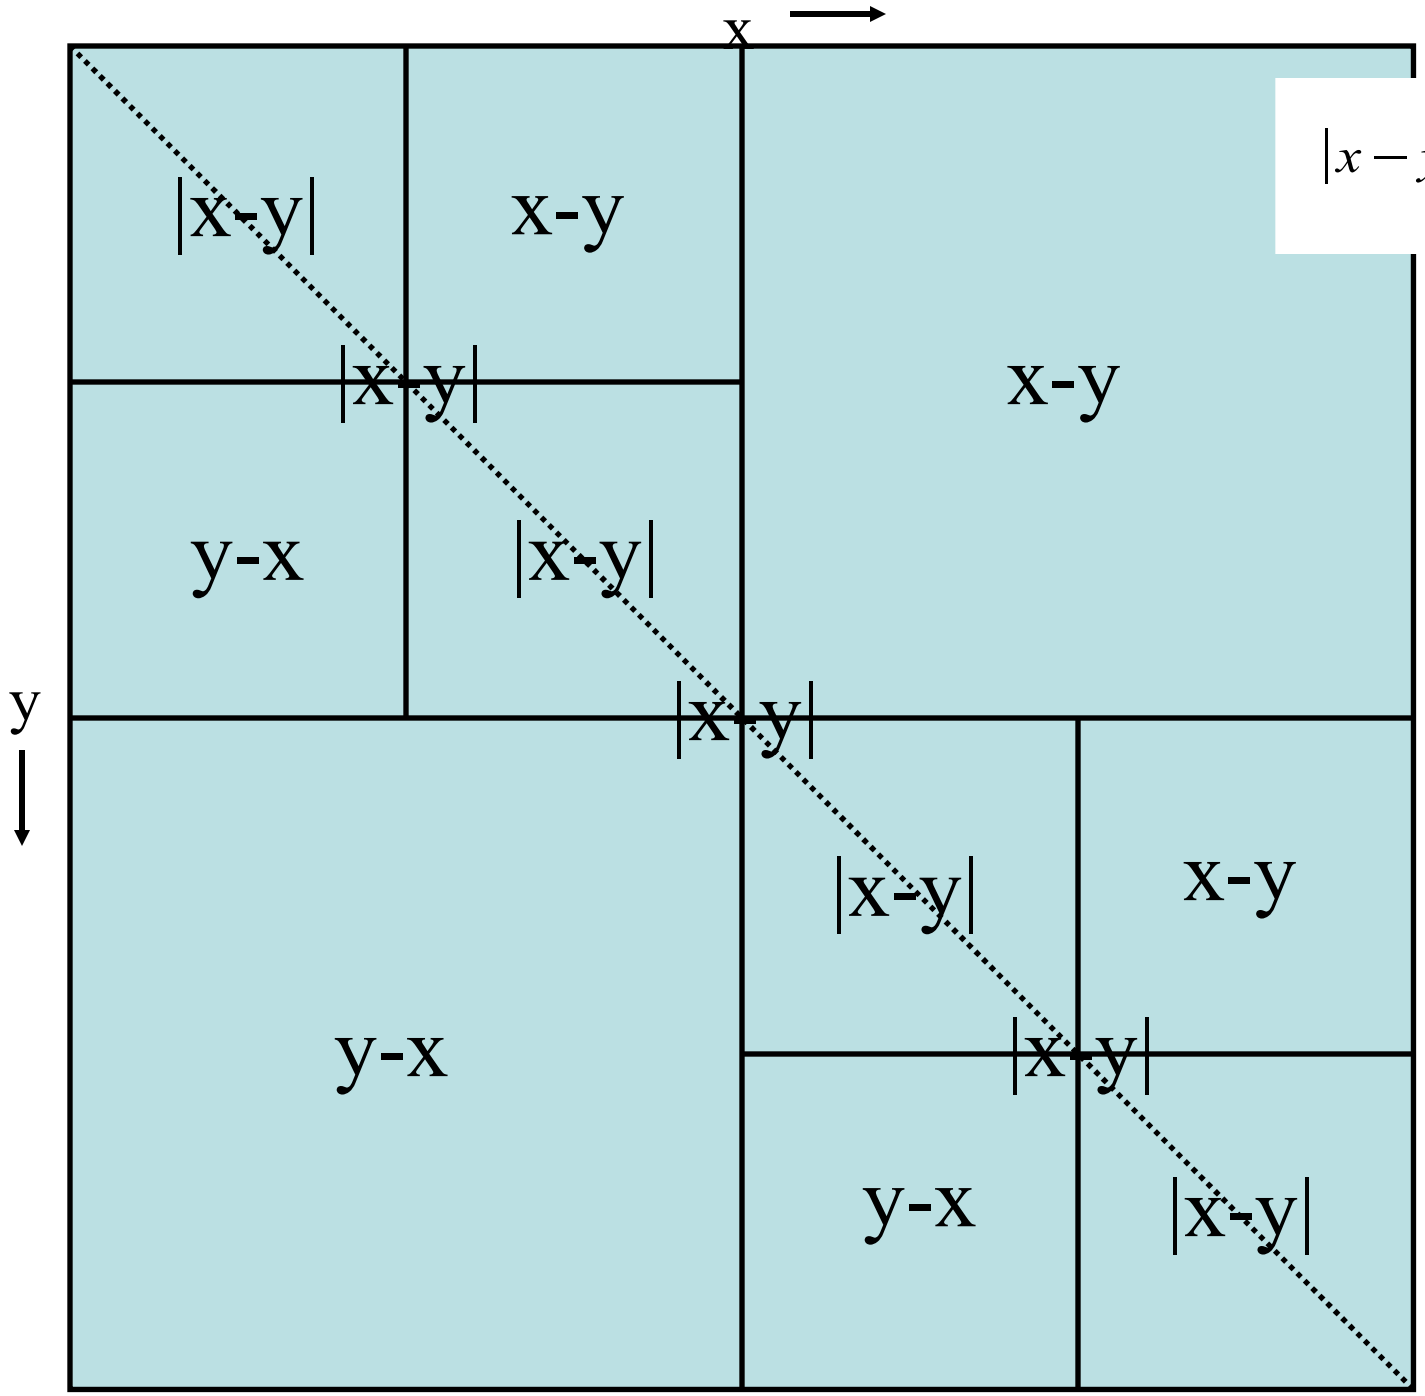

$$|x-y| = \sum_{\mu=1}^{r} f_\mu(x)\, g_\mu(y)$$

r = separation rank

In 3D, ideally must be one box removed from the diagonal

Diagonal box has full rank

Boxes touching diagonal (face, edge, or corner) have increasingly low rank

Away from diagonal $r = O(-\log \varepsilon)$

Integral Formulation

- Solving the integral equation
  - Eliminates the derivative operator and related “issues”
  - Converges as fixed point iteration *with no preconditioner*

$$\left(-\tfrac{1}{2}\nabla^2 + V\right)\Psi = E\Psi$$

$$\Psi = -2\left(-\nabla^2 - 2E\right)^{-1} V\Psi$$

$$= -2G*\left(V\Psi\right)$$

$$\left(G*f\right)(r) = \int ds \frac{e^{-k|r-s|}}{4\pi|r-s|} f(s) \quad \text{in 3D ; } k^2 = -2E$$

Such Green’s Functions (bound state Helmholtz, Poisson) can be rapidly and accurately applied with a single, sparse matrix vector product.

Separated form for integral operators

$$T * f = \int ds\, K(r-s) f(s)$$

- Approach
  - Represent the kernel over a finite range as a sum of products of 1-D operators (often, not always, Gaussian)

$$r^{n,l-l'}_{ii',jj',kk'} = \sum_{\mu=0}^{M} X^{n,l_x-l'_x}_{ii'} Y^{n,l_y-l'_y}_{jj'} Z^{n,l_z-l'_z}_{kk'} + O(\epsilon)$$

  - Only need compute 1D transition matrices (X,Y,Z)
  - SVD the 1-D operators (low rank away from singularity)
  - Apply most efficient choice of low/full rank 1-D operator
  - Even better algorithms not yet implemented

Accurate Quadratures

$$\frac{e^{-\mu r}}{r} = \frac{2}{\sqrt{\pi}} \int_0^\infty e^{-x^2 t^2 - \mu^2/4t^2}\, dt$$

$$= \frac{2}{\sqrt{\pi}} \int_{-\infty}^\infty e^{-x^2 e^{2s} - \mu^2 e^{-2s}/4 + s}\, ds$$

- Trapezoidal quadrature
  - Geometric precision for periodic functions with sufficient smoothness
- Beylkin & Monzon
  - Further reductions, but not automatic

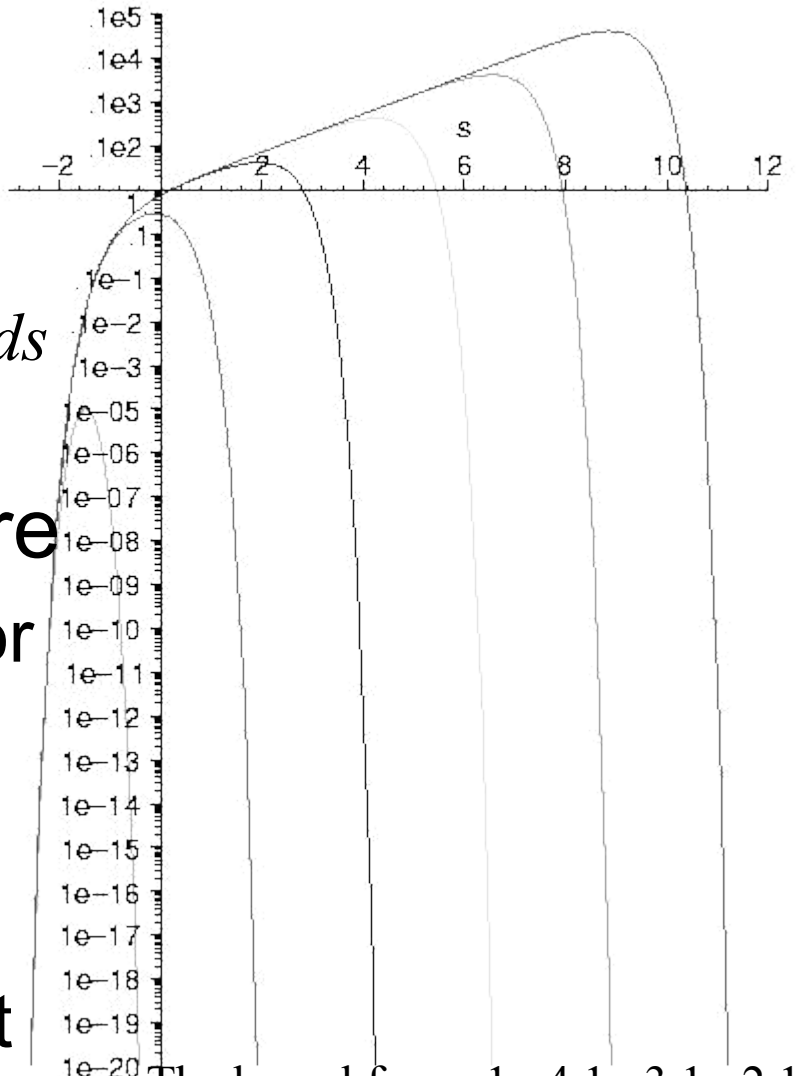

The kernel for x=1e-4,1e-3,1e-2,1e-,1e0. The curve for x=1e-4 is the rightmost

The Nature of Scattering Problems

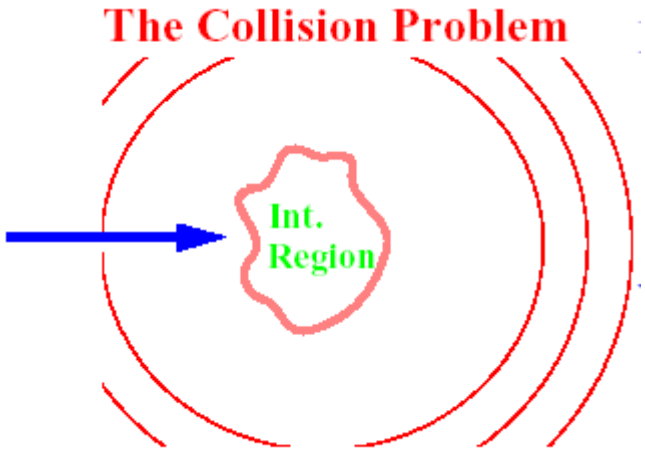

Map *known* "incoming" solutions onto *known* "outgoing" solutions

$$\Psi_{in} \Rightarrow \Psi_{interacting} \Rightarrow \Psi_{out}$$

Boundary conditions (e.g. one particle)

$$\Psi \rightarrow e^{i\mathbf{k}\bullet\mathbf{r}}_{in} + f(\vartheta,\varphi)e^{ikr}_{out}/r$$

Courtesy CW McCurdy

Why Are These Problems Difficult?

- E.g., double photoionization of atoms and molecules and electron-impact ionization are processes that place two electrons “in the continuum”

- The final state contains three separating charged particles

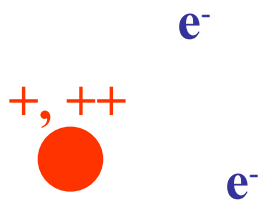

- All states, bound and continuum will be contained in the scattered wave.
- In the absence of correlation there would be essentially no cross section -- e.g., He:

$$\left\langle \varphi_{1s}(1)\varphi_{1s}(1)\left|\varepsilon\bullet r_1+\varepsilon\bullet r_2\right|\varphi_{ks}(1)\varphi_{kp}(2)\right\rangle \approx 0$$

Courtesy CW McCurdy

Time evolution

- Multiwavelet basis not optimal
  - Not strongly band limited
  - Explicit methods very unstable
    (DG introduces flux limiters, we use filters)
- Semi-group approach
  - Split into linear and non-linear parts

$$\dot{u}(x,t) = \hat{L}u + N(u,t)$$

$$u(x,t) = e^{\hat{L}t}u(x,0) + \int_0^t e^{\hat{L}(t-\tau)}N(u,\tau)\,d\tau$$

- Trotter-Suzuki methods
  - Time-ordered exponentials
  - Chin-Chen gradient correction (JCP 114, 7338, 2001)

$$e^{A+B} = e^{A/2}e^{B}e^{A/2} + O(\|[[A,B],A]\ldots\|)$$

Exponential propagator

- Imaginary time Schrodinger equation
  - Propagator is just the heat kernel

$$\left(-\frac{1}{2}\nabla^2+V(x)\right)\psi(x,t)=\dot{\psi}(x,t)$$

$$\psi(x,t)\simeq e^{\nabla^2 t/4}e^{-Vt}e^{\nabla^2 t/4}\psi(x,0)$$

$$e^{\nabla^2 t/2}f(x)=\frac{1}{\sqrt{2\pi t}}\int_{-\infty}^{\infty}e^{-\frac{(x-y)^2}{2t}}f(y)\,dy$$

$$\lim_{t\to\infty}\psi(x,t)=\psi_0(x)$$

- Wrap in solver to accelerate convergence

Exponential propagator

- Free-particle propagator in real time

$$\psi(x,t) = e^{i\nabla^2 t/2}\psi(x,0) = \frac{1}{\sqrt{2\pi i t}}\int_{-\infty}^{\infty} e^{-\frac{(x-y)^2}{2it}}\psi(y,0)\,dy$$

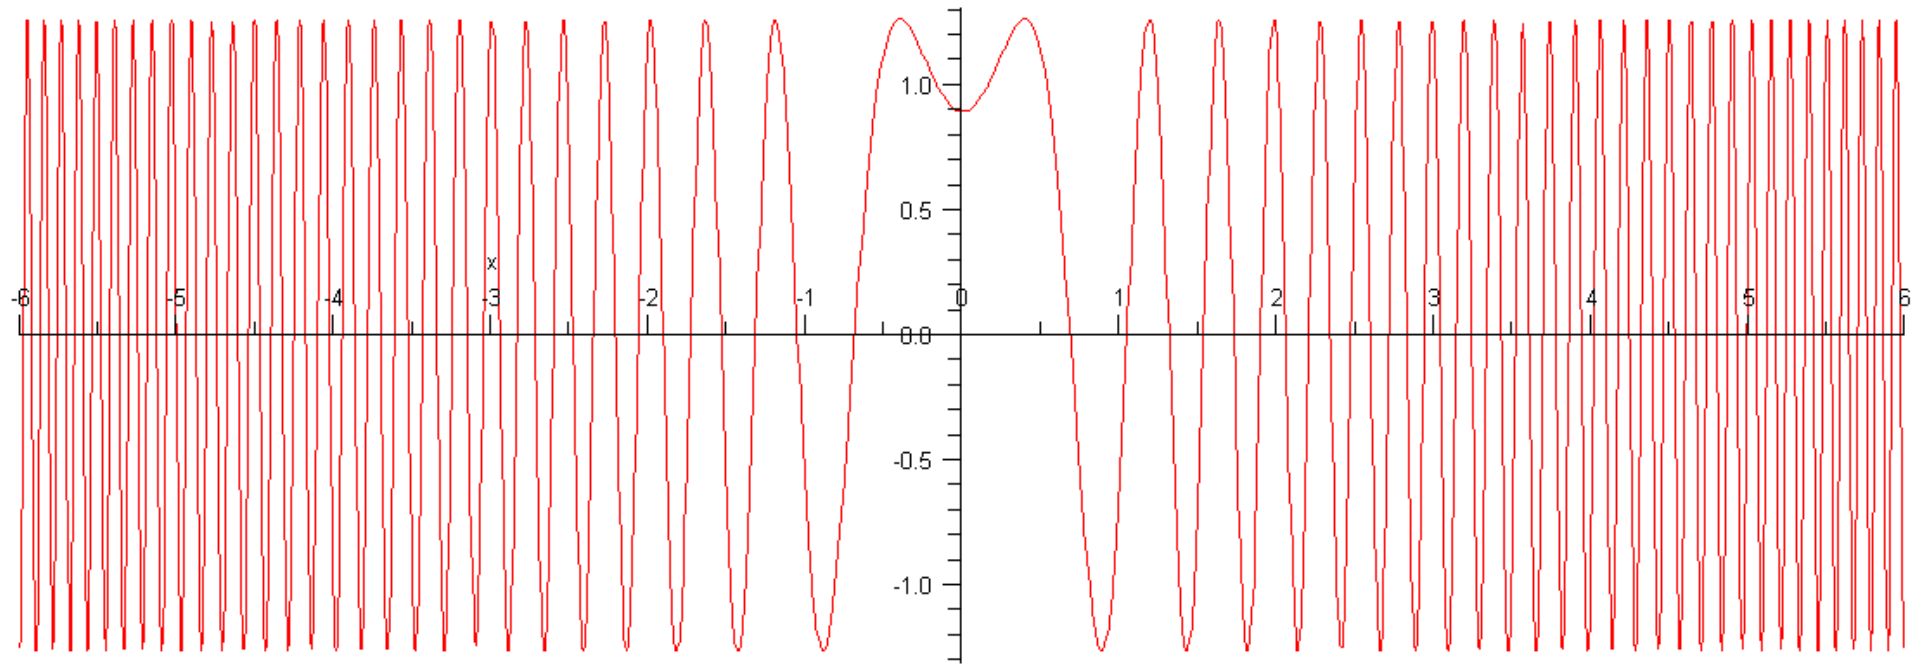

Exponential propagator

- Combine with projector onto band limit

$$\hat{G}_0(k,t,c) = e^{-i\frac{k^2 t}{2}}\left(1+(k/c)^{30}\right)^{-1}$$

$$h = \frac{\pi}{c} \qquad t_{crit} = \frac{2h^2}{pi}$$

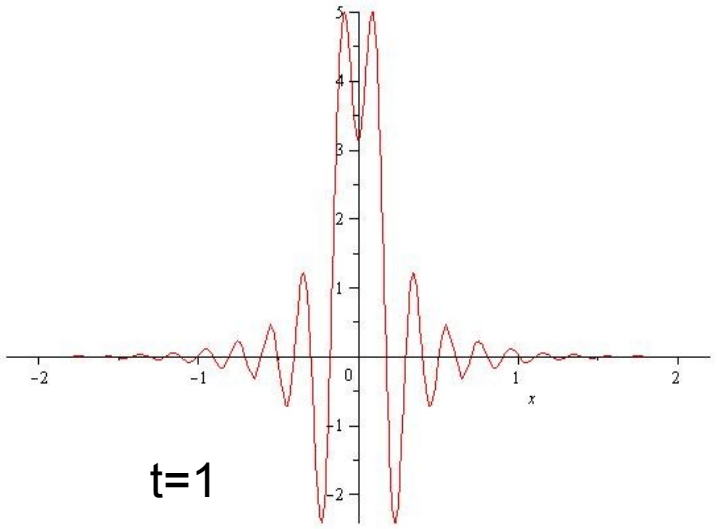

t=1

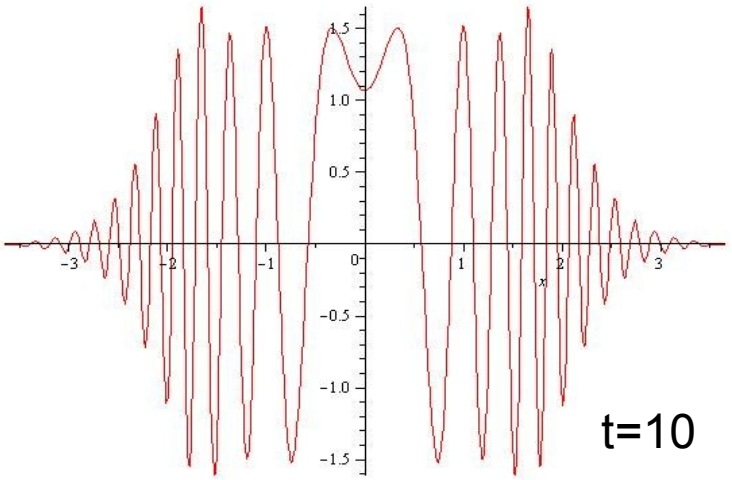

t=10

Path to linear scaling HF & DFT

- Need speed and precision
  - Absolute error cost $O\left(N \ln N / \epsilon\right)$
  - Relative error cost $O\left(N \ln 1 / \epsilon\right)$
- Coulomb potential
- HF exchange potential
- Orbital update
- Orthogonalization and or diagonalization
- Linear response properties

The Jaguar Cray XT4 Leadership System

2007

- 11,508 compute nodes
  - Dual-core AMD Opteron processors with 4 GB memory
  - 23,016 compute cores
- 396 service & I/O nodes
- ~750 TB local storage
- 3D Torus interconnect
- 46 TB aggregate memory
- 119 TF peak performance

2008

- 7,832 compute nodes
  - Quad-core AMD Opteron processors with 8 GB memory
  - 31,328 compute cores
- 240 service & I/O nodes
- 900 TB local storage
- 3D Torus interconnect
- 63 TB aggregate memory
- >250 TF peak performance
- General availability to user community in May 2008

ORNL Provides Leadership Computing to 2008 INCITE Program

- Allocation of computing resources to 30 programs in 2008 under the DOE's Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program (together with NERSC and PNNL).
- Leading researchers from government, industry, and the academic world will explore challenges including climate change, energy and alternative fuels on the center's leadership computers.
- This year's total allotment of processing hours nearly doubles that which ORNL provided in 2007.

Project Allocations: 145.3 million hrs

Industrial Allocations: 11.9 million hrs

New petaflop system will provide over 1B hours per year!

Some Science Drivers

Science Domains	Science and Engineering Driver
Accelerator Physics	Optimize a new low-loss cavity design for the ILC
Astrophysics	Explosion mechanism of core-collapse supernovae and Type Ia supernovae
Biology	Can efficient ethanol production offset the current oil and gasoline crisis?
Chemistry	Catalytic transformation of hydrocarbons; clean energy & hydrogen production and storage
Climate	Predict future climates based on scenarios of anthropogenic emissions
Combustion	Developing cleaner-burning, more efficient devices for combustion.
Fusion	Plasma turbulent fluctuations in ITER must be understood and controlled
High Energy Physics	Find the Higgs particles thought to be responsible for mass, and find evidence of supersymmetry
Nanoscience	Designing high temperature superconductors, magnetic nanoparticles for ultra high density storage
Nuclear Energy	Can all aspects of the nuclear fuel cycle be designed virtually? Reactor core, radio-chemical separations reprocessing, fuel rod performance, repository
Nuclear Physics	How are we going to describe nuclei whose fundamental properties we cannot measure?

ORNL INCITE 2008 Allocations by Discipline

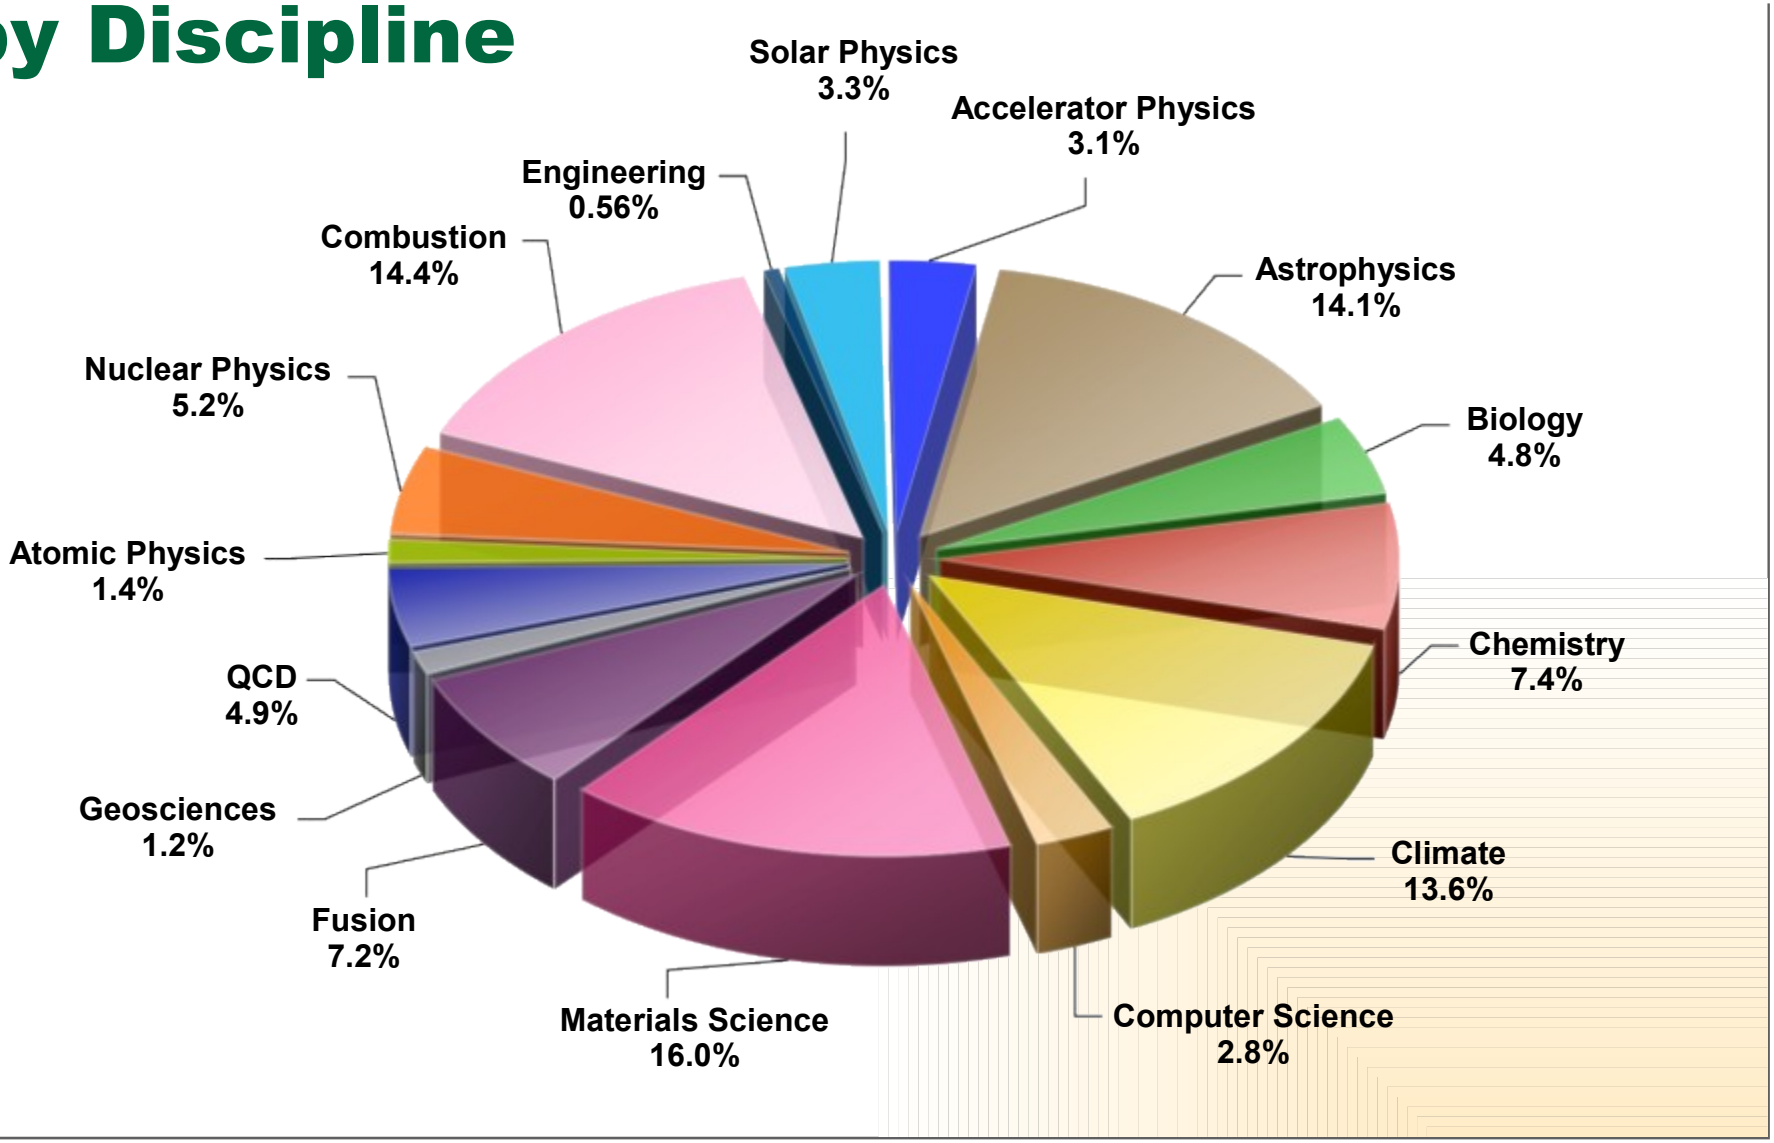

Our need for leadership computing

- Definitive, benchmark computations
  - The scale and fidelity we expect from petascale simulation will answer truly hard questions about real systems. Fully quantitative computations are central to fundamental understanding and to enabling rational design.
- Integration of experiment and theory
  - Fast turnaround of reliable simulations is already enabling the intimate integration of theory and simulation into chemistry, which is a predominantly experimental discipline.

The role of simulation in heavy element chemistry for advanced fuel cycles

- Molecular-scale knowledge is vital to enable the rational design of new/enhanced agents
  - Reduced cost & risk with increased efficiency
  - Current experimental approach can generate only a fraction of required data over many years
    - The rest are guesstimated.
  - We can compute much of this
    - Need higher precision than currently feasible
  - Combinatorial methods use thermodynamics for screening, but this is not reliable enough
- Approach
  - Mixed MM/QM Gibbs-free energy computations of partition coefficients
  - Simulation of select liquid-liquid, gas-gas interfaces
  - Accurate thermo-chemistry and spectroscopy
    - Many-body methods incorporating relativistic effects
- Outcomes
  - Design of new separation chemistries on a timescale relevant to engineering requirements (months to years rather than decades)

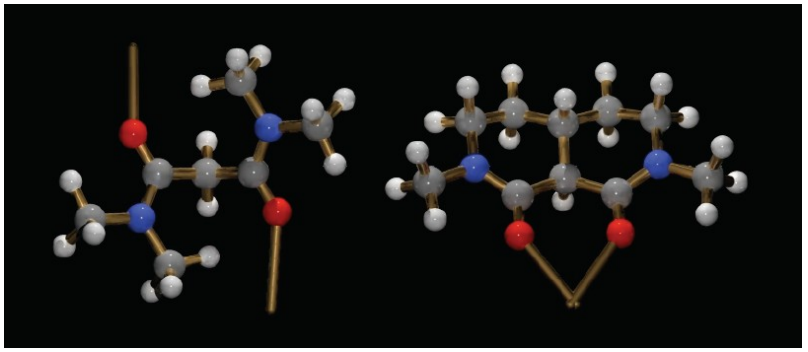

B. Hay, EMSP project 73759

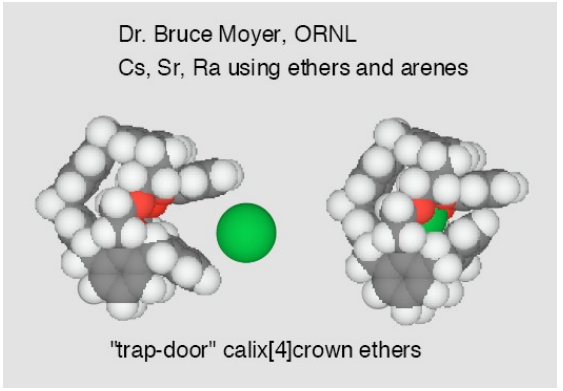

An Integrated Approach to the Rational Design of Chemical Catalysts

NCCS Incite Project

Robert J. Harrison: PI	**Oak Ridge National Laboratory, University of Tennessee**
Edoardo Aprà	**Oak Ridge National Laboratory**
Jerzy Bernholc	**North Carolina State University**
A.C. Buchanan III	**Oak Ridge National Laboratory**
Marco Buongiorno Nardelli	**North Carolina State University**
James M. Caruthers	**Purdue University**
W. Nicholas Delgass	**Purdue University**
David A. Dixon	**University of Alabama**
Sharon Hammes-Schiffer	**Pennsylvania State University**
Duane D. Johnson	**University of Illinois at Urbana Champaign**
Manos Mavrikakis	**University of Wisconsin at Madison**
Vincent Meunier	**Oak Ridge National Laboratory**
Mathew Neurock	**University of Virginia**
Steven H. Overbury	**Oak Ridge National Laboratory**
William F. Schneider	**University of Notre Dame**
William A. Shelton	**Oak Ridge National Laboratory**
David Sherrill	**Georgia Institute of Technology**
Bobby G. Sumpter	**Oak Ridge National Laboratory**
Kendall T. Thomson	**Purdue University**
Roberto Ansaloni	**Cray**
Carlo Cavazzoni	**CINECA, Italy**

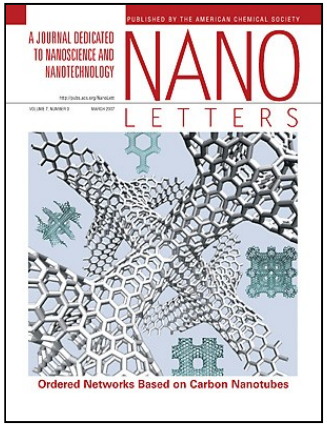

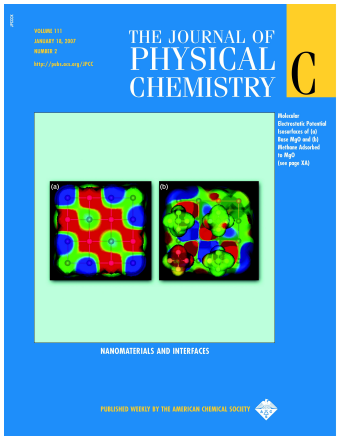

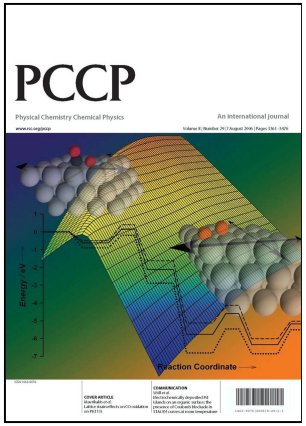

Amphoteric Doping of Carbon Nanotubes by Encapsulation of Organic Molecules: Electronic Transport and Quantum Conductance

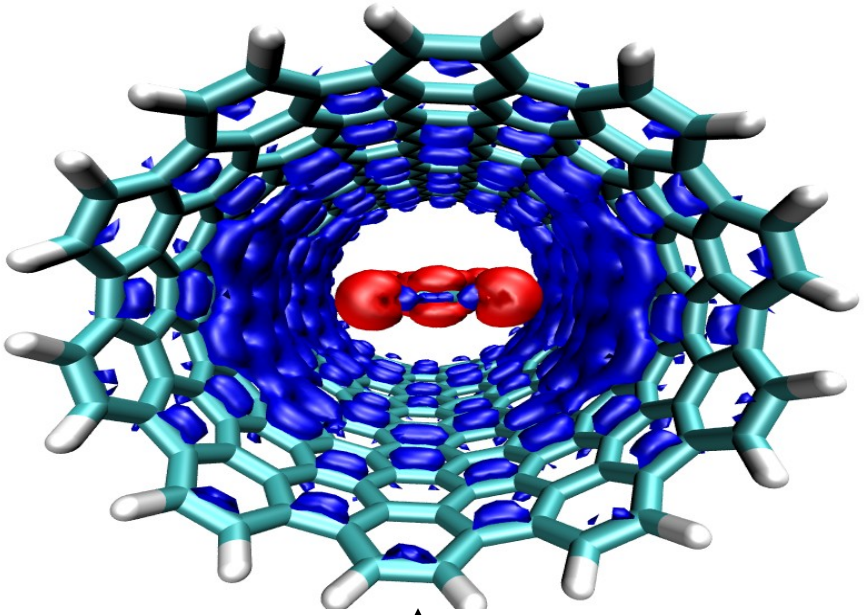

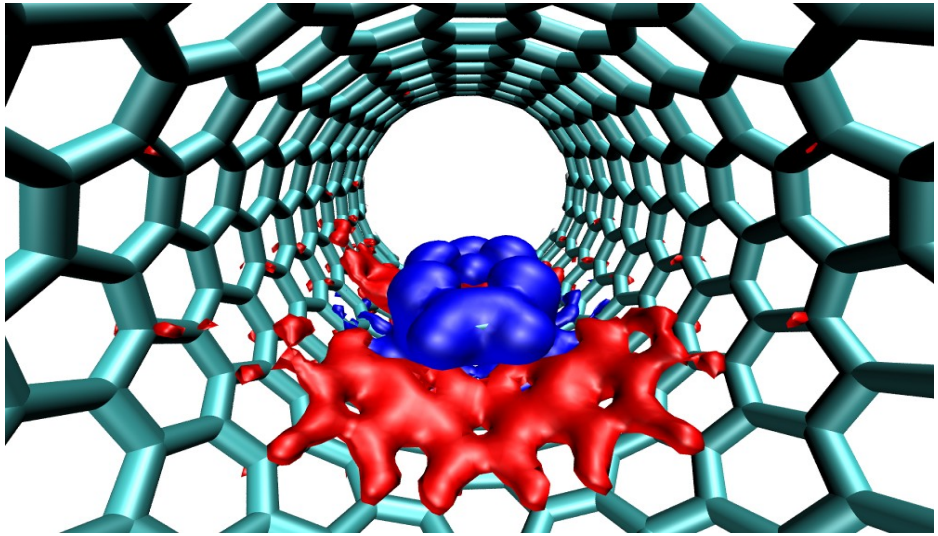

p-doped tube : holes are transferred from F_4-TCNQ to the nanotube

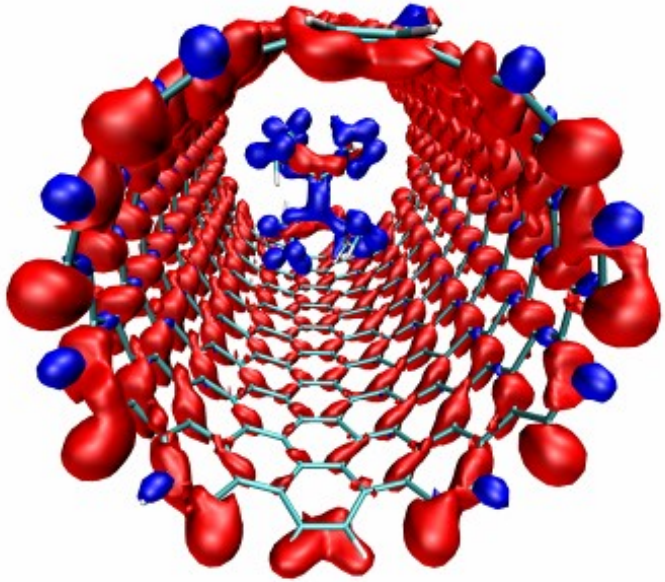

n-doped tube : holes are transferred from the tube to TTF and TDAE

Meunier, Sumpter

F4TCNQ: structure and transport

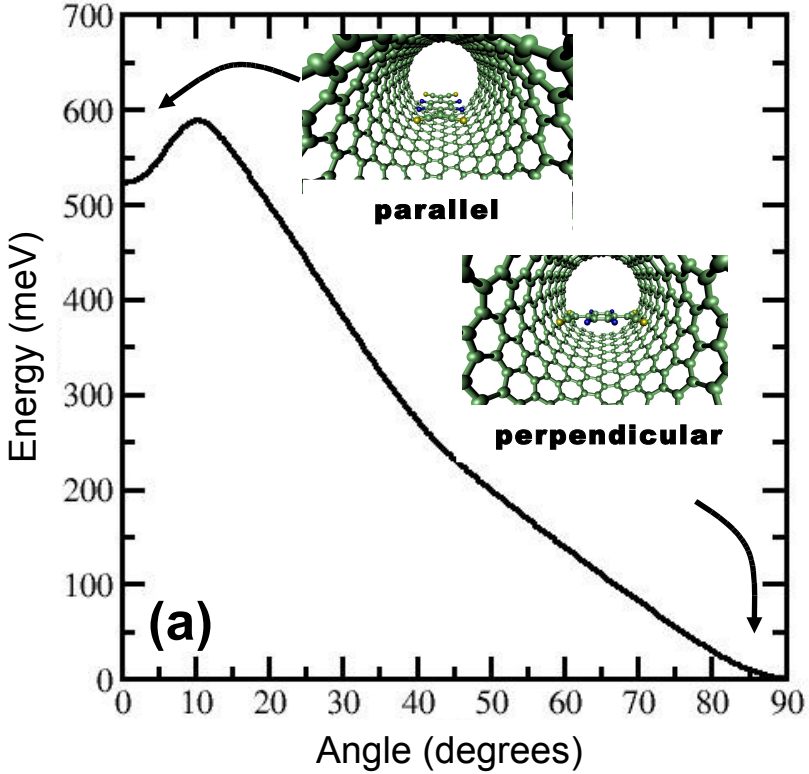

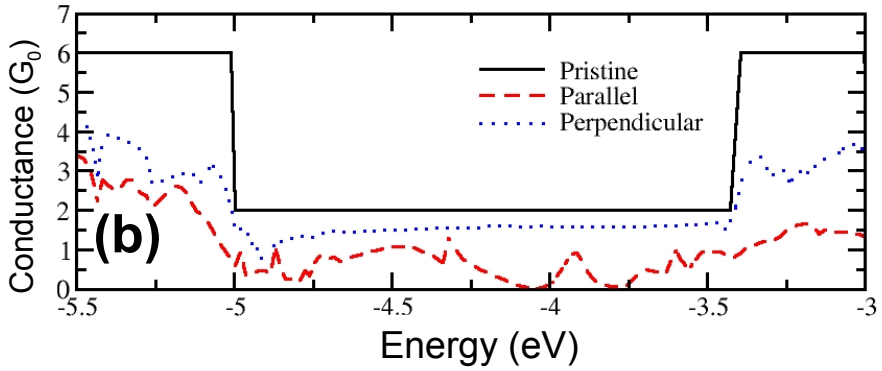

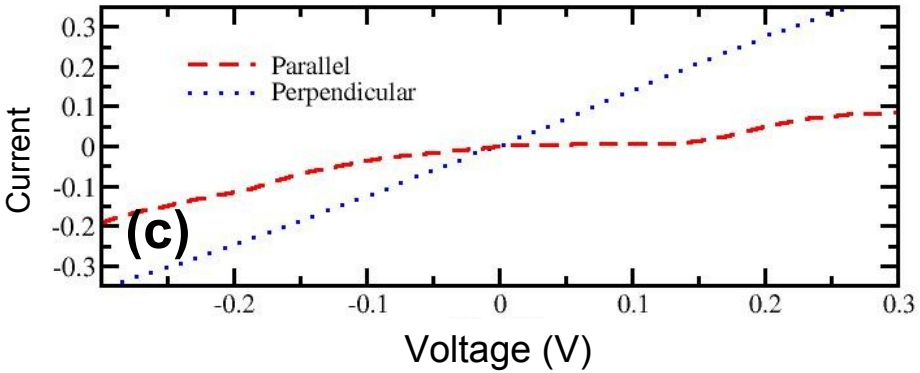

Kalinin, Meunier, Sumpter

Artificial DNA

Uses recent advances in synthetic biology for:

- Fundamental investigations to enable the development of sensor techniques for detecting Single Nucleotide Polymorphs

- Bionanowires and DNA by sequencing

- Probe the DNA replication process with unnatural DNA bases

- Force-field generation for the simulation of novel synthetic biological systems

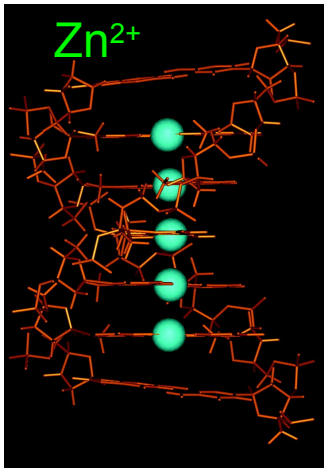

M-DNA

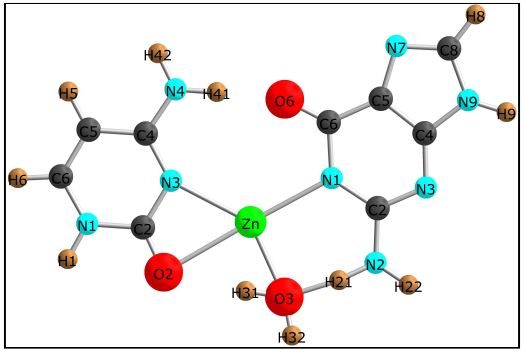

Metalated (M) and expanded (x) DNAs

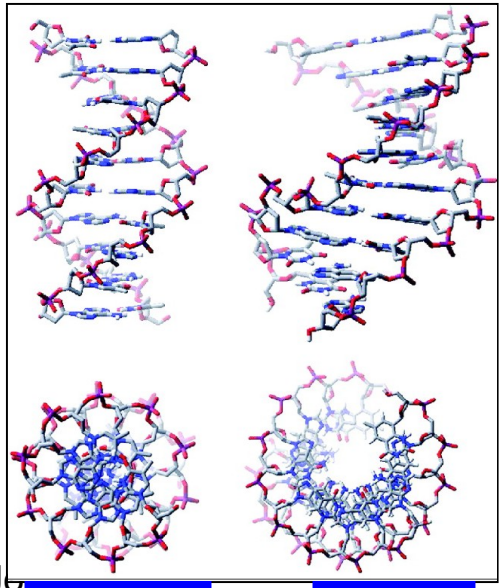

B-DNA xDNA

Fuentes, Šponer, Sumpter, Wells

Other technologies

- Field programmable gate arrays – multi TOP/s now

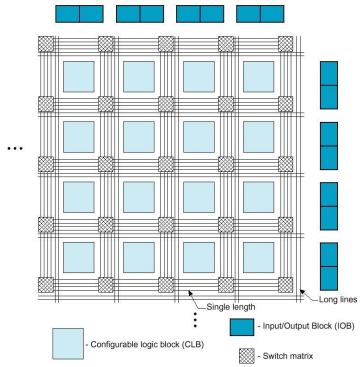

- General purpose graphical processor unit – 1TFLOP/s now
  - Lots of caveats on relevance

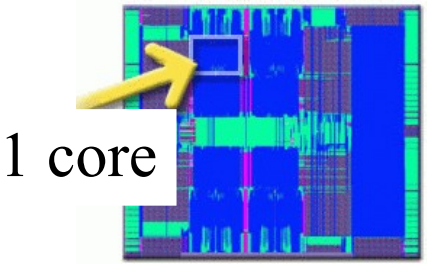

- Highly threaded devices
- FLOPs are cheap; bandwidth is expensive